Medical Decision Support: Experience with Implementing the Arden Syntax at the Columbia-Presbyterian Medical Center

Robert A. Jenders, MD, MS*†; George Hripcsak, MD*; Robert V. Sideli, MD*; William DuMouchel, PhD*‡; Haiying Zhang, MS‡; James J. Cimino, MD*†; Stephen B. Johnson, PhD*; Eric H. Sherman, MD, MS*†; Paul D. Clayton, PhD*†+

Departments of Medical Informatics*, Medicine†, Biostatistics‡ and Radiology+
Columbia University, New York, New York

We began implementation of a medical decision support system (MDSS) at the Columbia-Presbyterian Medical Center (CPMC) using the Arden Syntax in 1992. The Clinical Event Monitor which executes the Medical Logic Modules (MLMs) runs on a mainframe computer. Data are stored in a relational database and accessed via PL/I programs known as Data Access Modules (DAMs). Currently we have 18 clinical, 12 research and 10 administrative MLMs. On average, the clinical MLMs generate 50357 simple interpretations of laboratory data and 1080 alerts each month. The number of alerts actually read varies by subject of the MLM from 32.4% to 73.5%. Most simple interpretations are not read at all. A significant problem of MLMs is maintenance, and changes in laboratory testing and message output can impair MLM execution significantly. We are now using relational database technology and coded MLM output to study the process outcome of our MDSS.

INTRODUCTION

The integration of clinical guidelines into medical practice has been promoted by the development of computer-based reminder systems.[1] Strong evidence suggests that some medical decision support systems (MDSS) can improve physician performance.[2] The Arden Syntax for Medical Logic Modules (MLM), developed in part at Columbia-Presbyterian Medical Center (CPMC), has been promoted as an open standard for the procedural representation and sharing of medical knowledge.[3] The CPMC is a 1477-bed center. In 1994 it recorded 415256 outpatient clinic visits and 53374 discharges.

As part of our ongoing development of a MDSS at CPMC, we have been developing MLMs which serve the needs of clinical users, research scientists and hospital administrative personnel. We begin by discussing the architecture of the CPMC MDSS in which these MLMs reside and the software tools used to develop them. We then discuss the MLMs actually implemented at CPMC and the quantitative operation of those intended for clinical use. In reviewing our experience in this development process, we illustrate the benefits and problems associated with the implementation of our MDSS. Finally, we present an overview of our future work in studying the MDS and its impact on the process and outcome of medical care.

MDSS ARCHITECTURE

The MDSS at CPMC employs a compiler-interpreter pair to create executable code from the original syntax of the MLM.[4] Although a MLM can be created in this system using any text editor, we primarily employ a workstation with a sophisticated graphical user interface.[5] In particular, one of us has written an editor for MLMs employing Asterx running under X-Windows on a RS/6000 machine using Unix. MLMs so created are compiled into pseudocodes using a software program built with lex and yacc. These compiled routines then are uploaded to the main data repository of CPMC, a 3090 IBM mainframe, where they eventually execute.

Laboratory, demographic, admission-discharge-transfer (ADT) and inpatient pharmacy data are available for query by the MLMs. These data are stored in a relational format in DB2. Data access modules (DAMs) are PL/I programs which execute on the mainframe and allow the MLMs to query the database without requiring the MLM author to know the exact relational schema. Parameters passed from the read statement in a MLM to the DAM specify the type and identity of the data elements required and various constraints on the data such as inpatient status.

0195-4210/95/$5.00

Data stored in DB2 are encoded using the Medical Entities Dictionary developed at CPMC.[6] Each MLM can have one or more triggers which identify when the Clinical Event Monitor, which serves as the inference engine for the MDSS, should execute the MLM. These triggers are coded as pairs of MED codes, one of which identifies the type of data (such as a pharmacy order) and the other which signifies the identity of the actual data element whose storage in DB2 should fire the applicable MLM. The type of data element is specified in the trigger because the data model and the relational tables are partitioned in this manner.

In addition to storing patient data using MED codes, the output of a MLM can be written using MED codes. In most currently implemented MLMs, however, the interpretation, alert or electronic mail message is recorded in DB2 as a decision event whose value is the text string of the message. The application program which displays patient data for clinical users then queries the decision event tables to display appropriate information regarding a particular patient.

TYPES OF MLM

The MLMs currently implemented in our MDSS can be divided into three classes based on the intended audience of their output. The largest of these classes are the clinical MLMs. These include alerts which identify the presence of abnormal data or potentially serious combinations of data elements in the patient record. In addition, some clinical MLMs provide simple interpretations of laboratory data, such as calculating an anion gap or a calcium level corrected for albumin level from chemistry panel results. These interpretations do not provide any advice for further evaluation or treatment; they merely provide simple calculations derived from numbers in a test result.

The second most numerous class of MLMs is the research MLM. Generally these provide notice to clinical investigators via electronic mail when patients meeting certain characteristics are admitted to the hospital. In addition, these MLMs can notify researchers when a patient already enrolled in a study has been admitted to the hospital so that appropriate follow-up can occur.

The third principal class of MLMs are those intended for use by the hospital administrators, sometimes called Administrative Logic Modules (ALMs). These MLMs generally notify administrators via electronic mail when a particular condition has occurred so that corrective action can be taken in order to reduce hospital expenses. Examples of such MLMs include notice of inappropriate bed classification and live discharge from the hospital on the same day of admission. Also, this class of MLMs provides a service to affiliated community physicians by providing notice via fax or electronic mail of admissions or discharges of their patients from the hospital.

MLMs also can be classified by the kind of output they produce. Using this criterion, current MLMs can be classified as alert, interpretation, or electronic mail events. Most clinical MLMs are alerts, while most of the remaining MLMs produce output in the form of electronic mail. We are just now creating MLMs which produce output in the form of coded messages which can be used to represent intermediate decision states in a complicated clinical guideline.

IMPLEMENTED MLMS

We currently have 40 MLMs in production. Eighteen of these are clinical MLMs, including four interpretations and fourteen alerts (Table 1). The alert MLMs produce a summary average of 1080 distinct alerts each month, while the interpretation MLMs yield a summary average of 50357 events each month. Many more interpretations than alerts are produced because the former occur primarily on all valid chemistry panels, while the latter occur only for specific abnormal situations typically involving combinations of data elements.

Although eighteen distinct clinical MLMs exist, these may be grouped into several classes which act upon one kind of abnormal event. This occurs because distinct MLMs are created for triggers involving distinct kinds of events. For example, three MLMs are used to alert users to the presence of hypokalemia and digoxin use which might lead to cardiac dysrhythmia. One MLM is triggered by the storage of a pharmacy order for digoxin; another MLM is triggered by the storage of a blood potassium result; and the third MLM in the suite is triggered by the storage of a blood digoxin level. Thus, three different triggers with three different queries are used to encode this example, although similar logic is used in each of the MLMs.

SUBJECT	TYPE	NUMBER OF MLMs	NUMBER OF MONTHS	TOTAL EVENTS	AVERAGE EVENTS PER MONTH	STANDARD DEV
positive TB culture	alert	1	32	1098	34.3	11.9
creatinine clearance	interpretation	1	17	69689	4099.4	3925.9
fractional excretion of sodium	interpretation	1	25	3110	124.4	50.25
newborn with mother HBV positive	alert	2	36	450	12.5	7.6
hypokalemia and digoxin use	alert	3	37	3224	87.1	23.9
creatinine rise	alert	1	30	7623	254.1	126.8
chem 7	interpretation	1	21	415724	19796.4	5378.6
chem 20	interpretation	1	21	553084	26337.3	5084.8
hypokalemia and diuretic use	alert	2	3	198	66	16.6
renal failure and aminoglycoside use	alert	2	3	168	56	20.7
new anemia	alert	1	3	1292	430.7	289.7
renal failure and NSAID use	alert	2	3	417	139	96.0

Table 1. Operation statistics on clinical MLMs. STANDARD DEV is the standard deviation. TB is infection with tuberculosis. NSAIDS are nonsteroidal anti-inflammatory drugs. HBV is Hepatitis B Virus. Chem 7 and chem 20 are 7- and 20-element chemistry panels, respectively.

In addition to these clinical MLMs, we currently have twelve research MLMs in production. Examples of these include identifying patients with abnormal cervical pathology, patients with a MB fraction of creatine kinase positive for myocardial infarction, patients with a discharge diagnosis of lymphosarcoma and patients with hyperamylasemia suggestive of pancreatitis. These notify the appropriate researchers of the medical record number and, when appropriate, the inpatient location of a patient in order to provide the researchers with the opportunity to enroll the patient in a study. Other MLMs also notify researchers when a study patient is admitted to the hospital.

Finally our MDSS includes ten administrative MLMs. Examples of these are notifying community physicians when their patients are admitted and discharged, notifying administrators of live discharge from the intensive care unit to outpatient status, and identification of patients whose payment status is self-pay.

USER INTERACTION WITH DECISION EVENTS

Clinical alerts for each patient may be displayed on the hospital information system at the option of the individual user while looking at a particular patient's data. This option is available as a choice on a menu

SUBJECT	NUMBER OF EVENTS	NUMBER OF EVENTS VIEWED	NUMBER OF VIEWINGS	AVERAGE VIEWINGS PER EVENT	PERCENT OF EVENTS VIEWED
positive TB culture	257	189	1416	7.5	73.5
creatinine clearance	60889	1121	1436	1.3	1.8
fractional excretion of sodium	1663	530	993	1.9	31.9
newborn with mother HBV positive	209	115	219	1.9	55.0
hypokalemia and digoxin use	1202	691	1425	2.1	57.5
creatinine rise	1549	977	3048	3.1	63.1
chem 7	233334	3017	3391	1.1	1.3
chem 20	361992	4847	5309	1.1	1.3
hypokalemia and diuretic use	198	95	154	1.6	48.0
renal failure and aminoglycoside use	168	70	109	1.6	41.7
new anemia	1292	419	864	2.1	32.4
renal failure and NSAID use	417	145	262	1.8	34.8

Table 2. User interaction with clinical MLMs. These data concern only those events which occurred after the institution of logging of users viewing interpretations and alerts. Average Viewings Per Event applies only to those events (interpretations or alerts) which were viewed by at least one user.

which also indicates the date of the most recent alert. If at least one alert for a given patient has not been viewed by that particular user, a special message appears informing the user of that fact.

If the user chooses the menu option to view the alerts, he is then presented with a list of the titles of the alerts in reverse chronological order. The user then may choose an individual alert from the list in order to read its text. Also, interpretations of certain laboratory tests may be viewed by choosing an option while reading the result of a test for which an interpretation is available.

Data are available on which user read the text of which alert on which patient since January, 1994. These data are summarized in Table 2. The total number of events are those alerts and interpretations which occurred after the initiation of the logging of displayed alerts. The number of alerts or interpretations whose text actually was viewed and the percentage of the total number of such events that these viewings constitute also are indicated. In addition, the average number of viewings for those alerts actually viewed is delineated.

DISCUSSION

The relatively small percentage of viewings of particular MLM subjects may be explained by the manner in which alerts are displayed. If users choose the option to display alerts while viewing a patient's results, they first are given a list of alerts for that patient consisting only of the title of the alert. The title alone may be a sufficient alert for some users, who then do not read the text of the alert. The data in this paper do not include viewings of the title of an alert without reading its text.

Also, maintenance is a key issue in providing clinical and research system users with reliable service. One event which highlighted the importance of maintenance was the change in laboratory computer systems in Summer, 1994. This caused the introduction of a large number of new MED codes for the new laboratory tests encoded by the new laboratory software. As a result, the MLMs

then in existence would work with the few legacy laboratory tests encoded with old MED codes but often failed to function with newly generated laboratory data. Clinical users and researchers, who had become accustomed to the alerts, interpretations and notifications provided by the MDSS noted their absence. This problem helped to produce the relatively high standard deviations seen with the clinical MLMs.

This problem was solved by making use of class information encoded in the MED. In the MED hierarchy, many laboratory tests are grouped into classes which identify related kinds of tests. Instead of using the code of specific tests in the MLMs, we adapted them to employ the code of the class. The DAMs were adapted further to allow query by class. This allows future laboratory tests to be added to the relevant classes without disrupting the operation of the MLMs.

Another problem with maintenance occurred when the laboratory without prior notification changed the format of certain textual results. MLMs which looked for specific strings no longer functioned properly. A statistical monitoring system notified MLM developers that certain MLMs had stopped functioning, and a review of the new laboratory output allowed repair of the relevant MLMs.

FUTURE WORK

All of the MLMs used in production to date have produced output in the form of decision events whose value are unconstrained text strings which then are displayed to the user of the clinical information system either via an alert, an interpretation or electronic mail. In order to represent the intermediate decision steps of a complex clinical guideline, we are developing MLMs which write decision events whose values are MED codes. This enables MLMs to execute at an indeterminate future time using the already-calculated result of other MLMs. Writing coded decision values allows as-yet unspecified MLMs to execute using previously produced decisions at an indeterminate future time point.

In addition we currently are studying the impact of alerts on length of stay and various clinical outcomes.

SUMMARY

The CPMC MDSS includes 40 MLMs which execute on a mainframe computer and access patient data stored in a relational database via PL/I-encoded DAMs. Eighteen clinical MLMs provide a large number of alerts and data interpretation, while research and administrative MLMs provide study enrollment screening and warnings regarding cost to these different classes of users. Maintenance is a key issue in developing the MDSS, and using the MED to trigger and query by class has helped ease this process. We now are using coded MLM output to help study the impact of our MLMs.

References

1. Jenders RA, Morgan M, Barnett GO. Use of open standards to implement health maintenance guidelines in a clinical workstation. Comput Biol Med 1994; 24:385-390.

2. Johnston ME, Langton KB, Haynes RB, Mathieu A. Effects of computer-based clinical decision support systems on clinician performance and patient outcome. Ann Intern Med 1994; 120:135-142.

3. Hripcsak G, Clayton PD, Pryor TA, Haug P, Wigertz OB, Van der lei J. The Arden Syntax for medical logic modules. Editor Miller RA. Proceedings of the Fourteenth Annual Symposium on Computer Applications in Medical Care. New York: IEEE Computer Press, 1990; 200-204.

4. Hripcsak G, Cimino JJ, Johnson SB, Clayton PD. The Columbia-Presbyterian Medical Center decision-support system as a model for implementing the Arden Syntax. Proceedings of the Fifteenth Annual Symposium on Computer Applications in Medical Care. Editor Clayton PD. New York: McGraw-Hill, 1992; 248-252.

5. Hripcsak G, Clayton PD, Cimino JJ, Johnson SB, Friedman C. Medical decision support at Columbia-Presbyterian Medical Center. In Software Engineering in Medical Informatics. Editors Timmers T, Blum BI. Amsterdam: Elsevier Science 1991; 471-479.

6. Cimino JJ, Clayton PD. Coping with changing controlled vocabularies. Proceedings of the Eighteenth Annual Symposium on Computer Applications in Medical Care. Editor Ozbolt JG. Philadelphia: Hanley & Belfus, 1994; 135-139.